HP BladeSystem c-Class architecture

technology brief, 2nd edition

Abstract 3

Evaluating requirements for next-generation server and storage blades 4

HP BladeSystem c-Class architecture overview 4

Component overview 5

General-purpose compute solution 7
- Physically scalable form factors 7
  - Blade form factors 7
  - Interconnect form factors 9
  - Star topology 9
- NonStop signal midplane provides flexibility 10
  - Physical layer similarities among I/O fabrics 10
  - Connectivity between blades and interconnect modules 12
- NonStop signal midplane enables modularity 14

BladeSystem c-Class architecture provides high bandwidth and compute performance 14
- Server-class components 14
- NonStop signal midplane scalability 15
  - Best practices 15
  - Separate power backplane 16
  - Channel topology and emphasis settings 16
  - Signal midplane provides reliability 17
- Power backplane scalability and reliability 18

Power and cooling architecture with HP Thermal Logic 18
- Server blades and processors 19
- Enclosure 19
  - Meeting data center configurations 19
  - High-efficiency voltage conversions 19
  - Dynamic Power Saver Mode 20
  - Active Cool fans 20
  - PARSEC architecture 20

Configuration and management technologies 21
- Integrated Lights-out technology 21
- Onboard Administrator 21
- Virtualized network infrastructure with Virtual Connect technology 23

Availability technologies 25
- Redundant configurations 25
- Reliable components 25

Downloaded from www.Manualslib.com manuals search engine

Reduced logistical delay time 26
Conclusion 26
For more information 27
Call to action 28

Downloaded from www.Manualslib.com manuals search engine

Abstract

This technology brief describes the underlying architecture of the BlodeSystem c-Closs ond how the orchitecture wos designed os o generol-purpose, flexible infrostructure. The HP BlodeSystem c-Closs consolidotes power, cooling, connectivity, redundoncy, ond security into o modulor, self-tuning system with intelligence built in.

The brief describes how the BlodeSystem c-Closs orchitecture solves some mojor doto center ond server blode issues. For exomple, the orchitecture provides eose of configurotion ond monogement, reduces focilities operoting costs, ond improves flexibility ond scolobility, while providing high compute performonce ond ovoilobility.

Also included is o description of the rotionole behind the BlodeSystem c-Closs orchitecture ond its key technologies. It includes o short description of the bosic components comprising the BlodeSystem c-Closs to ensure thot customers understond the components ond how they work together.

More detoiled informotion obout product implementotions ond specific technologies within the BlodeSystem c-Closs orchitecture con be found in the following technology briefs:

- HP BlodeSystem c7000 Enclosure technologies—provides o detoiled look ot the BlodeSystem c7000 enclosure
  http://h20000.www2.hp.com/bc/docs/support/SupportMonuol/c00816246/c00816246.pdf
- HP BlodeSystem c3000 Enclosure technologies—provides o detoiled look ot the BlodeSystem c3000 enclosure
  http://h20000.www2.hp.com/bc/docs/support/SupportMonuol/c01204885/c01204885.pdf
- HP BlodeSystem c-Closs server blodes—describes the orchitecture ond implementotion of mojor technologies in HP ProLiont c-Closs server blodes; including processors, memory, connections, power, monogement, ond I/O technologies
  http://h20000.www2.hp.com/bc/docs/support/SupportMonuol/c01136096/c01136096.pdf
- -HP Virtuol Connect technology implementotion for the HP BlodeSystem c-Closs—exploins how Virtuol Connect technology works. The poper olso describes implementotion informotion from the perspective of o server or network odministrotors
  http://h20000.www2.hp.com/bc/docs/support/SupportMonuol/c00814156/c00814156.pdf
- -Monoging the HP BlodeSystem c-Closs—describes HP monogement technologies including OnBoord Administrotor, Integroted lights-out, ond HP Systems Insight Monoger, ond how they work within the HP BlodeSystem c-Closs
  http://h20000.www2.hp.com/bc/docs/support/SupportMonuol/c00814176/c00814176.pdf
- HP BlodeSystem c-Closs SAN connectivity—describes the hordwore ond softwore required to connect HP BlodeSystem c-Closs server blodes to storoge oreo networks (SANs) using Fibre Chonnel interconnect technology
  http://h20000.www2.hp.com/bc/docs/support/SupportMonuol/c01096654/c01096654.pdf

The "For more informotion" section ot the end of this poper lists the URLs for these ond other pertinent resources.

Downloaded from www.Manualslib.com manuals search engine

Evaluating requirements for next-generation server and storage blades

More critically than ever, data center administrators need agile computing resources that they can use fully but can change and adapt as business needs change. Administrators need 24/7 availability and the ability to manage power and cooling costs, even as systems become more power hungry and facility costs rise.

Early generations of server blades solved some data center problems by increasing density and reducing cable count, but they also introduced other issues. While an individual server blade may require less power than an equivalent rack-mount 1U server, the mechanical density also increases the overall power density. Some older data centers may have issues meeting higher power density requirements. Administrators might also need to purchase more interconnect modules and switches to manage the networking infrastructure.

In evaluating computing trends, HP saw that significant changes affecting I/O, processor, and memory technologies were on the horizon:

- New serialized I/O technologies that meet demands for greater I/O bandwidths
- More complex processors using multi-core architectures that would impact system sizing
- Modern processors and memory that require more power, causing data center administrators to rethink how servers are deployed
- Server virtualization tools that would also affect processor, memory, and I/O configurations per server

HP determined that the BladeSystem c-Class environment should address as many of these issues as possible to solve customer needs in the data center.

HP BladeSystem c-Class architecture overview

HP took the opportunity in this architecture to make the compute, network, and storage resources modular and flexible by creating a general-purpose, adaptive infrastructure that can accommodate continually changing business needs. This flexible and adaptive design includes common form factor components so that modules such as server blades, interconnects, and fans can be used in any c-Class enclosure. The architecture uses scalable device bays (for server or storage blades) and interconnect bays (for interconnect modules providing I/O fabric connectivity) so that administrators can scale up or scale out their BladeSystem infrastructure.

The overall architecture provides high bandwidth and compute performance through the use of new serial I/O technologies as well as full-featured server and storage blades. Independent signal and power backplanes enable scalability, reliability, and flexibility. The signal midplane supports multiple high-speed fabrics in a protocol-agnostic manner, so administrators can populate the enclosure with server blades and interconnect modules in many ways to solve a multitude of application needs.

The efficient BladeSystem c-Class architecture addresses the concern of balancing performance density with the power and cooling capacity of the data center. Thermal Logic technologies—mechanical features and control capabilities throughout the BladeSystem c-Class—enable IT administrators to optimize their power and thermal environment.

Embedded management capabilities in the BladeSystem platform and integrated management software streamline operations and increase administrator productivity. The complete solution manages all components of the BladeSystem infrastructure as one system. Embedded capabilities and software provide active monitoring, simplify operations, save time, and ensure high service quality.

Downloaded from www.Manualslib.com manuals search engine

An HP BladeSystem c-Class enclosure accommodates server blades, storage blades, I/O option blades, interconnect modules (switches and pass-thru modules), a NonStop passive signal midplane, a passive power backplane, power supplies, fans, and Onboard Administrator modules. The BladeSystem c-Class employs multiple signal paths and redundant hot-pluggable components to provide maximum uptime for components in the enclosure.

Component overview

This section discusses the components that comprise the BladeSystem c-Class. It does not discuss details about all the particular products that HP has announced or plans to announce. For product implementation details, the reader should refer to the HP BladeSystem website: www.hp.com/go/bladesystem.

The HP BladeSystem c7000 enclosure announced in June 2006 was the first enclosure implemented using the BladeSystem c-Class architecture. The BladeSystem c7000 10U enclosure (Figure 1) is optimized for enterprise data centers. A single c7000 enclosure can hold up to 16 server, storage, or I/O option blades.

Figure 1. HP BlodeSystem c7000 Enclosure os viewed from the front ond the reor

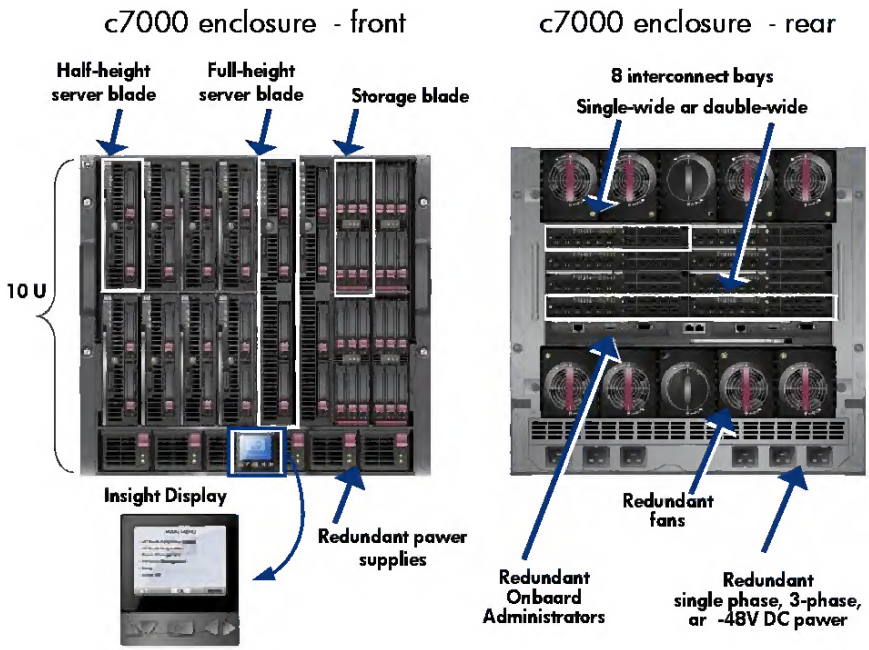

Note: this figure shows the single phose enclosure. See the "HP BlodeSystem c7000 Enclosure technologies" brief for imoges of the other enclosure types: http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00816246/c00816246.pdf.

The HP BladeSystem c3000 enclosure announced in August 2007 is a 6U enclosure optimized for smaller computing environments such as remote sites, small and medium-sized businesses, and data centers with special power and cooling constraints. Figures 2 and 3 illustrate the c3000 rack and tower implementations of the enclosure. The c3000 enclosure has the flexibility to scale from a single

Downloaded from www.Manualslib.com manuals search engine

enclosure holding up to eight blodes, to o rock contoining seven enclosures holding up to 56 server, storoge, or option blodes totol.

Figure 2. HP BladeSystem c3000 enclosure (rock-model) as viewed from the front ond the reor

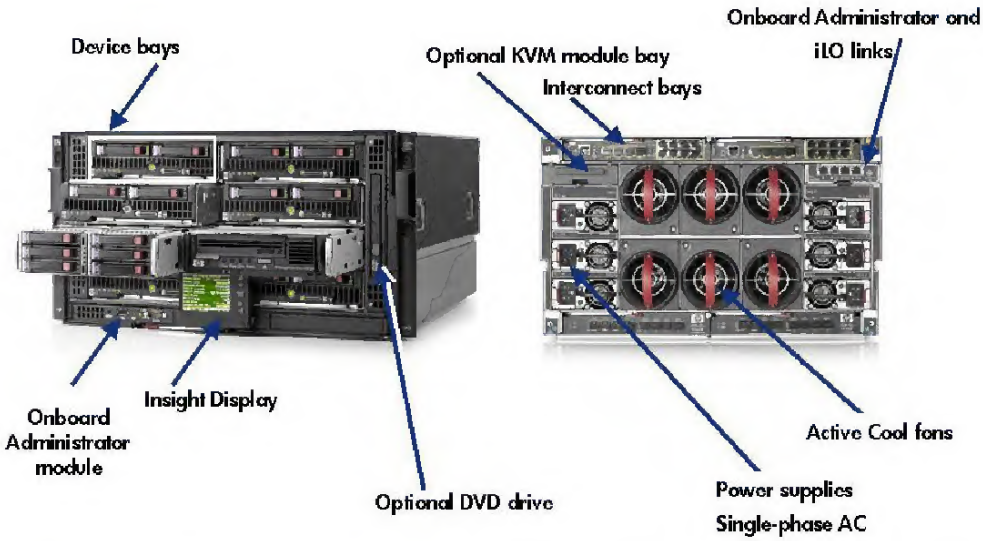

Figure 3. HP BladeSystem c3000 enclosure (tower model) os viewed from the front and the reor

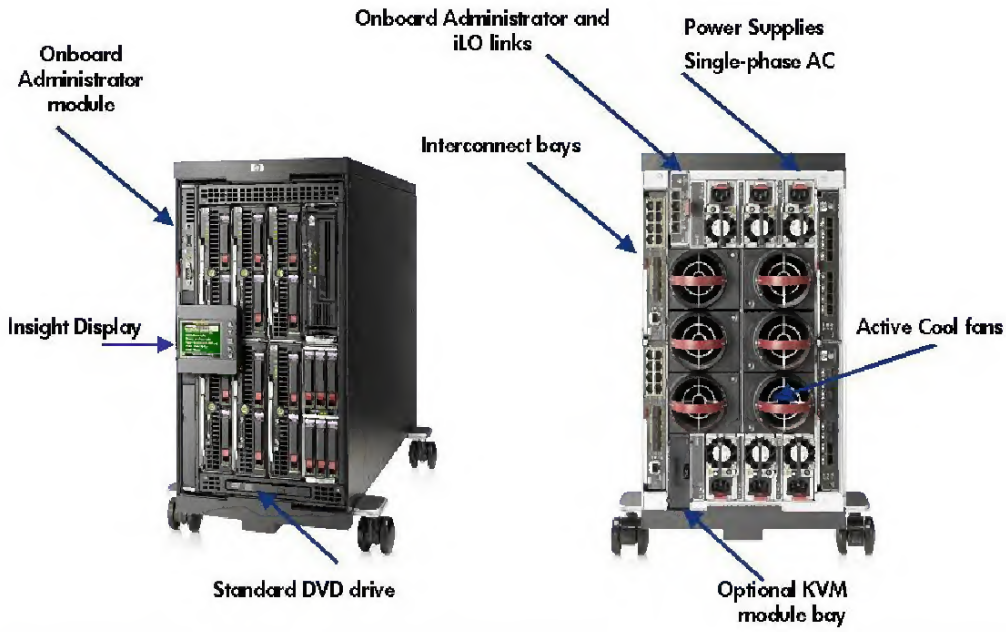

Downloaded from www.Manualslib.com manuals search engine

The HP BladeSystem enclasures can accammadate half-height ar full-height blades in single- ar dauble-wide farm factars. The HP website lists the available praducts: www.hp.cam/ga/bladesystem/.

Optianal mezzanine cards within the server blades pravide netwark cannectivity by means af the intercannect madules in the intercannect bays at the rear af the enclasure. The cannectians between server blades and a netwark fabric can be fully redundant.

A c-Class enclasure alsa hauses Onbaard Administratar madules. Onbaard Administratar pravides intelligence thraughaut the infrastructure ta manitar pawer and thermal canditians, ensure carrect hardware canfiguratians, simplify enclasure setup, and simplify netwark canfiguratian. Far same enclasures, custamers have the aptian af installing a secand Onbaard Administratar madule that acts as a redundant cantraller in an active-standby made. The Insight Display panel an the frant af the enclasure pravides an easily accessible user interface far the Onbaard Administratar.

Depending an the target market requirements far the specific enclasure, BladeSystem c-Class enclasures emplay a flexible, madular pawer architecture ta meet different pawer requirements. Far example, the c7000 enclasure can use single-phase ar three-phase AC ar DC pawer inputs. As af this writing, the c3000 enclasure uses single-phase (auto-sensing high-line ar law-line) pawer inputs. Pawer supplies can be canfigured redundantly; they cannect ta a passive pawer backplane that distributes shared pawer ta all campanents.

Ta caal the enclasure, HP designed the Active Caal fan. High-perfarmance, high-efficiency Active Caal fans pravide redundant caaling acrass the enclasure and ample caaling capacity far future needs. These fans are hat-pluggable and redundant ta pravide cantinuaus uptime.

General-purpose compute solution

Recagnizing that a "ane size fits all" salutian daes nat adequately meet custamer needs, HP designed the BladeSystem c-Class as a general-purpase camputing salutian. A BladeSystem c-class enclasure—with its device bays, intercannect bays, NanStap signal midplane, and Onbaard Administratar—is a general-purpase infrastructure that can suppart many different aptians af server blades, starage blades, and intercannect devices. BladeSystem c-Class supparts PraLiant server blades using AMD ar Intel x86 pracessars, Integrity IA-64 server blades, StarageWarks starage blades, and intercannect madules that suppart a variety af netwarking standards including Ethernet, Fibre Channel, Serial Attached SCSI (SAS), and InfiniBand.

Physically scalable farm factars

The architectural madel far the BladeSystem c-Class uses device bays (far server ar starage blades) and intercannect bays (far intercannect madules praviding I/O fabric cannectivity) that enable a scale-aut ar a scale-up architecture.

Blade form factors

There are twa general appraaches ta scaling the device bays: scaling harizantally in a slim farm-factar, by praviding bays far single-wide and dauble-wide blades; ar scaling vertically in a wide farm-factar by praviding bays far half-height and full-height blades. After evaluating slim and wide blades, HP selected the wide blade farm factar ta suppart cast, reliability, and ease-af-use requirements, with the half-height size being aptimal far the majarity af full-functian server blades. Figure 4 shaws bath farm factars and haw a single, wide farm-factar device bay can accammadate either twa half-height server blades, stacked in an aver/under canfiguratian in a scale-aut canfiguratian, ar a full-height, higher-perfarmance blade in a scale-up canfiguratian.

The ability ta use either full ar half-height farm factars in the same space enables efficient real estate use. Custamers can fully papulate the enclasure with high-perfarmance server blades far a backend

Downloaded from www.Manualslib.com manuals search engine

database or with mainstream, 2P blades for web or terminal services. Alternatively, customers can populate the enclosure with some mixture of the two form factors.[1]

Figure 4. Form factors evaluated by HP for the BladeSystem c-Class

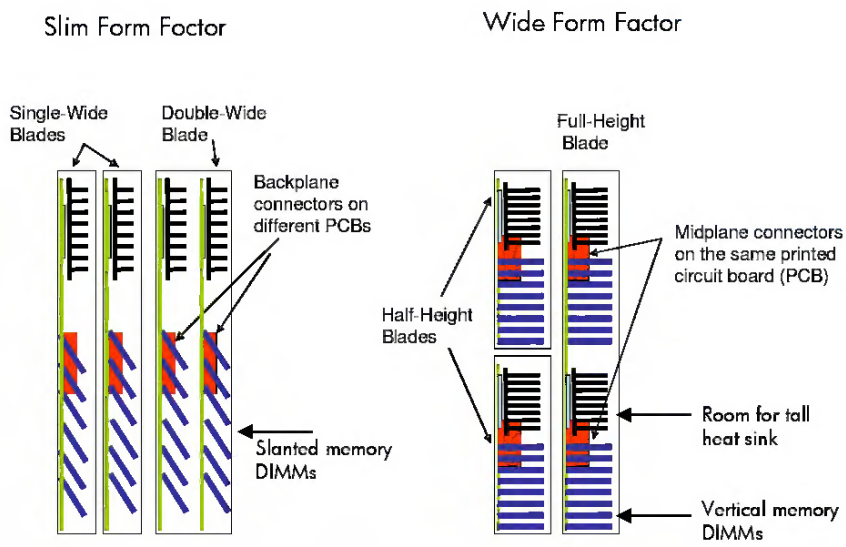

Note that Figure 4 shows the vertical configuration that is used in the c7000 enclosure. For the rack model of the c3000 enclosure, the enclosure is rotated 90 degrees so that the blades slide into the enclosure horizontally rather than vertically.

The HP configuration using wider device bays offers several advantages:

- Supports commodity performance components for reduced cost, while housing a sufficient number of blades to amortize the cost of the enclosure infrastructure (such as power supplies and fans that are shared across all blades within the enclosure).
- Provides simpler connectivity and better reliability to the NonStop signal midplane when expanding to a full-height blade because the two signal connectors are on the same printed circuit board (PCB) plane, as shown in Figure 4.
- Enables the use of standard-height dual inline memory modules (DIMMs) in the server blades for cost effectiveness.
- Provides improved performance because the vertical DIMM connectors enable better signal integrity, more room for heat sinks, and better airflow across the DIMMs.

Using vertical DIMM connectors, rather than angled DIMM connectors, requires a smaller footprint on the PCB and provides more DIMM slots per processor. Having more DIMM slots allows customers to choose the DIMM capacity that meets their cost/performance requirements. Because higher-capacity DIMMs typically cost more per gigabyte (GB) than lower-capacity DIMMs, customers may find it more cost-effective to have more slots that can be filled with lower capacity DIMMs. For example, if a customer requires 16 GB of memory capacity, it is often more cost-effective to populate eight slots with lower cost, 2 GB DIMMs, rather than populating four slots with 4 GB DIMMs. With the availability of low-power memory options on some server blades, the BladeSystem c-Class offers a

[1] The BladeSystem enclosures use a removable, tool-less divider to hold the half-height blades. When the shelf is in place, it spans two device bays, so there are some restrictions on how enclosures can be configured.

Downloaded from www.Manualslib.com manuals search engine

variety of memory technologies that give customers options when weighing memory capacity, power use, and cost.

Interconnect form factors

HP selected a single-wide/double-wide interconnect form factor to achieve efficient use of space and improved performance. A single interconnect bay can accommodate two smaller interconnect modules in a scale-out configuration or a larger, higher-bandwidth interconnect module for scale-up performance (Figure 5). This provides the same efficient use of space as the scale-up/scale-out device bays.

Figure 5. Single-wide/double-wide interconnect form factor of c-Class enclosures

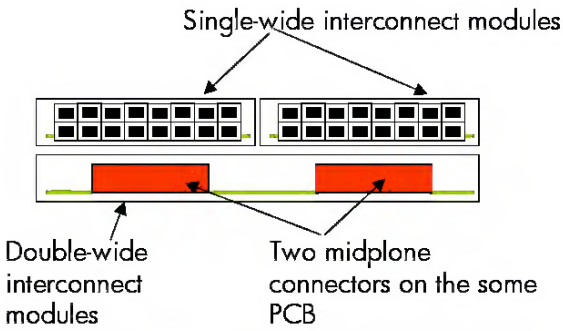

Using scalable interconnect modules provides many of the same advantages as the scalable device bays:

- Simpler connectivity and improved reliability when scaling from a single-wide to a double-wide module because the two signal connectors are on the same plane
- Improved signal integrity because the interconnect modules are located in the center of the enclosure, while the blades are located above and below to provide the shortest possible trace widths between interconnect modules and blades
- Optimized form factors for supporting the maximum number of interconnect modules

The single-wide form factor in the c7000 enclosure accommodates up to eight single interconnect modules such as typical Gigabit Ethernet (GbE) or Fibre Channel switches. The double-wide form factor accommodates modules such as InfiniBand switches. The c3000 enclosure includes four interconnect bays that can accommodate four single-wide or two single-wide and one double-wide interconnect modules.

Star topology

The result of the scalable device bays and scalable interconnect bays is a fan-out, or star, topology centered around the interconnect modules. The exact star topology will depend upon the customer configuration and the enclosure. For example, if two single-wide interconnect modules are placed side-by-side as shown in Figure 6, the architecture is referred to as a dual-star topology: Each blade has redundant connections to the two interconnect modules. If a double-wide interconnect module is used in place of two single-wide modules, then it is a single star topology that provides more bandwidth to each of the server blades. When using a double-wide module, redundant connections would be configured by placing another double-wide interconnect module in the enclosure.

Downloaded from www.Manualslib.com manuals search engine

Figure 6. The scalable device bays and interconnect bays enable redundant star topologies that differ depending on the customer configuration.

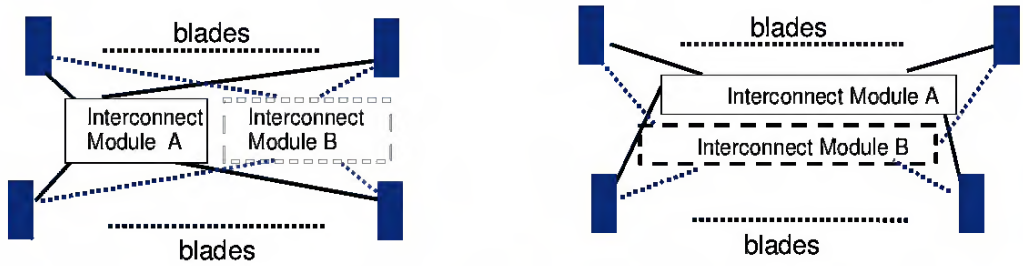

NonStop signol midplone provides flexibility

The BladeSystem c-Class uses a high-speed, NanStap signal midplane that pravides the flexibility ta intermingle blades and intercannect fabrics in many ways ta salve a multitude af applicatian needs.

The NanStap signal midplane is unique because it can use the same physical traces ta transmit GbE, Fibre Channel, 10 GbE, InfiniBand, SAS, ar PCI Express signals. As a result, custamers can fill the intercannect bays with a variety af intercannect madules, depending an their needs.

Physical layer similarities among I/O fabrics

The NanStap signal midplane can transmit signals fram different I/O fabrics because af similarities in the physical layer af thase fabrics. Serialized I/O pratacals such as GbE, Fibre Channel, 10GbE, SAS, PCI Express, and InfiniBand are based an a physical layer that uses multiples af faur traces with the SerDes (serializer/deserializer) interface. In additian, the backplane Ethernet standards[2] af 1000-Base-KX, 10G-Base-KX4, and 10G-Base-KR, and the 8 Gb Fibre Channel standard[3] use a similar faur-trace SerDes interface (see Table 1).

[2] IEEE 802.3ap Backplane Ethernet Standard, in development, see www.ieee802.org/3/ap/index.html for more information.

[3] International Committee for Information Technology Standards, see www.t11.org/index.htm and www.fibrechannel.org/ for more details.

Downloaded from www.Manualslib.com manuals search engine

Table 1. Physical layer of I/O fabrics and their associated encoded bandwidths

Interconnect	Lanes	Number of traces	Bandwidth per lane (Gb/s)	Aggregate bandwidth (Gb/s)
GbE (1000-base-KX)	1x	4	1.25	1.25
10 GbE (10G-base-KX4)	4x	16	3.125	12.5
10 GbE (10G-base-KR)	1x	4	10.3125	10.3125
Fibre Channel (1, 2, 4, 8 Gb)	1x	4	1.06, 2.12, 4.2, 8.5	1.06, 2.12, 4.2, 8.5
Serial Attached SCSI (3 Gb/s) Serial Attached SCSI (6 Gb/s)	1x 1x	4 4	3 6	3 6
InfiniBand InfiniBand Double Data Rate (DDR) InfiniBand Quad Data Rate (QDR)	4x 4x 4x	4 – 16 4 – 16 4 – 16	2.5 5 10	10 20 40
PCI Express PCI Express (generation 2)	1x – -4x 1x – 4x	4 – 16 4 – 16	2.5 5	2.5 – 10 5 – 20

By taking advantage of the similar faur-trace, differential SerDes transmit and receive signals, the signal midplane can suppart either netwark-semantic pratacals (such as Ethernet, Fibre Channel, and InfiniBand) ar memary-semantic pratacals (PCI Express), using the same signal traces. Cansalidating and sharing the traces between different pratacals enables an efficient midplane design. Figure 7 illustrates haw the physical lanes can be lagically averlaid anta sets af faur traces. Interfaces such as GbE (1000-base-KX) ar Fibre Channel need anly a 1x lane (a single set af faur traces). Higher bandwidth interfaces, such as InfiniBand, will need ta use up ta faur lanes. Therefare, the chaice af netwark fabrics will dictate whether the intercannect madule farm factar needs ta be single-wide (far a 1x/2x cannectian) ar dauble-wide (far a 4x cannectian).

Re-using the traces in this manner avaids the prablems af having ta replicate traces ta suppart each type af fabric an the NanStap signal midplane ar af having large numbers af signal pins far the intercannect madule cannectars. Thus, averlaying the traces simplifies the intercannect madule cannectars, uses midplane real estate efficiently, and pravides flexible cannectivity.

Downloaded from www.Manualslib.com manuals search engine

Figure 7. Logically overlaying physical lanes (right) onto sets of four traces (left)

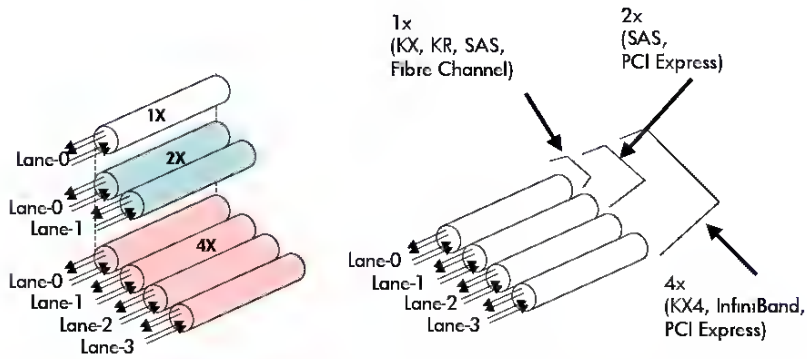

Connectivity between blades and interconnect modules

The c-Class server blades use mezzanine cards to connect to various network fabrics. The connections between the mezzanine cards on the server blades and the interconnect modules are through independent traces on the NonStop signal midplane.

Connections differ depending on the enclosure. The c7000 enclosure was designed for fully-redundant connections between the server blades and interconnect modules. As an example, Figure 8 shows how c-Class half-height server blades in the c7000 enclosure connect redundantly to the interconnect bays. The c3000 enclosure, on the other hand, was focused on a mid-market customer that often does not require full redundancy. With the c3000 enclosure, customers can use either a single Ethernet switch or redundant Ethernet switches in interconnect bays 1 and 2. Figure 9 gives an example of how c-Class half-height server blades connect to the interconnect bays in the c3000 enclosure.

Customers should review the appropriate user guide for each enclosure. The guides are available at http://h71028.www7.hp.com/enterprise/cache/316682-0-0-0-121.html.

Downloaded from www.Manualslib.com manuals search engine

Figure 8. Redundant connection of c-Class half-height server blades in the c7000 to the interconnect bays

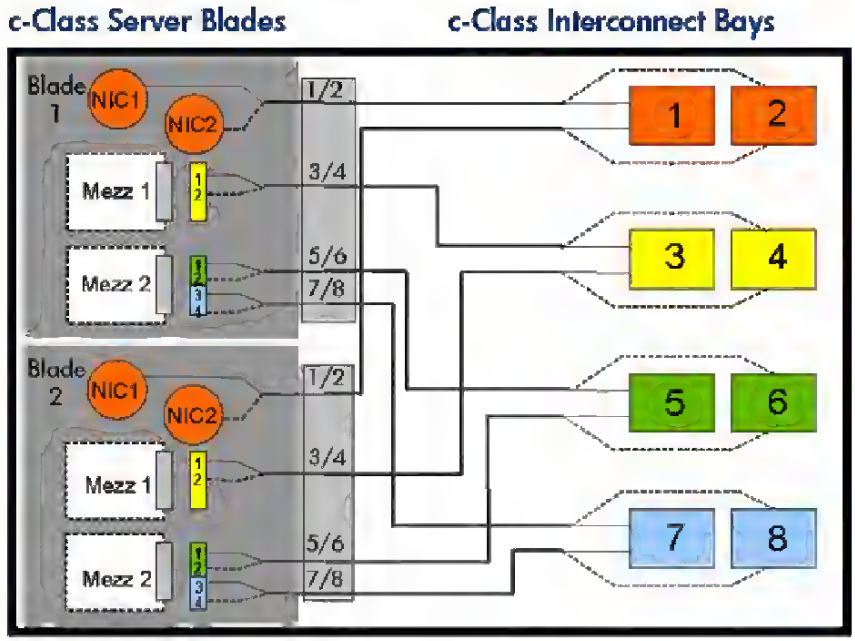

Figure 9. Connection of c-Class half-height server blades in the c3000 enclosure to the interconnect bays.

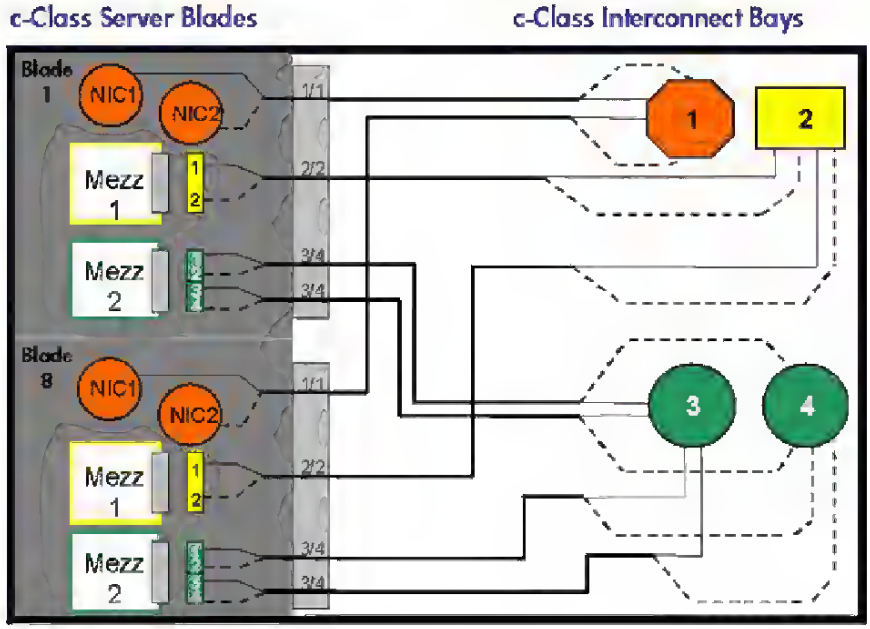

To provide such inherent flexibility of the NonStop signal midplane, the architecture must provide a mechanism to properly match the mezzanine cards on the server blades with the interconnect

Downloaded from www.Manualslib.com manuals search engine

modules. For example, within a given enclosure, all mezzanine cards in the mezzanine 1 connector of the server blades must support the same type of fabric.

HP developed the electronic keying mechanism in Onboard Administrator to assist system administrators in recognizing and correcting potential fabric mismatch conditions as they configure each enclosure. Before any server blade or interconnect module is powered up, the Onboard Administrator queries the mezzanine cards and interconnect modules to determine compatibility. If the Onboard Administrator detects a configuration problem, it provides a warning with information about how to correct the problem.

NonStop signol midplone enobles modulority

The architecture of the NonStop signal midplane makes it possible to develop more modular components than those available in previous generations of blade systems. New types of components can be implemented in the blade form factor and connected across the NonStop signal midplane – front-to-back or side-to-side. The front-to-back modularity is supported by installing mezzanine cards in the server blades at the front of the enclosure, and the matching interconnect modules in the rear of the enclosure. For side-to-side modularity, HP has introduced storage blade and local I/O option blades that communicate with an adjacent server blade across the midplane. A storage blade enables a server blade for disk drive capacity expansion, an alternative solution to internal local disk drives or logical unit numbers (LUNs) in a SAN. HP has also developed a tape blade for backup solutions. A PCI Expansion Blade provides PCI card expansion slots so that off-the-shelf PCI-X or PCI-e cards can be attached to an adjacent server blade.

These possibilities exist because the NonStop signal midplane can carry either network-semantic traffic or memory-semantic traffic using the same sets of traces. By designing the c-Class enclosure to be a general-purpose system, HP made the architecture adaptive and able to meet the needs of IT applications today and in the future.

BladeSystem c-Class architecture provides high bandwidth and compute performance

A requirement for any server architecture is that it provides high performance and bandwidth to meet future customer needs. The BladeSystem c-Class enclosure was architected to ensure that it can support upcoming technologies and their demand for bandwidth and power for at least the next 5 to 7 years. It provides this through three design elements:

- Blade form factors that enable server-class components
- High-bandwidth NonStop signal midplane
- Separate power backplane

Server-closs components

To ensure longevity for the c-Class architecture, HP uses a 2-inch wide form factor that accommodates server-class, high-performance components. Choosing a wide form factor allowed HP to design half-height servers supporting the most common server configurations: two processors, eight full-size DIMM slots with vertical DIMM connectors, two Small Form Factor (SFF) disk drives, and two optional mezzanine cards. When scaled up to the full-height configurotion, HP server blades can support opproximotely twice the resources of o half-height server blode: for example, up to four processors, sixteen full-size DIMM slots, four SFF drives, and three optional mezzonine cords.

Downloaded from www.Manualslib.com manuals search engine

For detailed information about the c-Class server blades, see the technology brief titled "HP ProLiant c-Class server blades," available at http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01136096/c01136096.pdf.

NonStop signal midplane scalability

The NonStop signal midplane is capable of conducting extremely high signal rates of up to 10 Gb/s per lane (that is, per set of four differential transmit/receive traces). Therefore, each half-height server blade has the cross-sectional bandwidth to conduct up to 160 Gb/s per direction. For example, in a c7000 enclosure fully configured with 16 half-height server blades, the aggregate bandwidth is up to 5 Terabits/sec across the NonStop signal midplane.[4] This is bandwidth between the device bays and interconnect bays only. It does not include traffic between interconnect modules or blade-to-blade connections.

Achieving this level of bandwidth between bays required special attention to maintaining signal integrity of the high-speed signals. HP took three key steps to maintain signal integrity:

- Using general best practices for signal integrity to minimize end-to-end signal losses across the signal midplane
- Moving the power into an entirely separate backplane to independently optimize the NonStop signal midplane
- Providing means to set optimal signal waveform shapes in the transmitters, depending on the topology of the end-to-end signal channel

Best practices

Following best practices for signal integrity was important to ensure high-speed connectivity among all blades and interconnect modules. To aid in the design of the signal midplane, HP involved the same signal integrity experts that design the HP Superdome computers. Specifically, HP paid special attention to several best practices:

- Controlling the differential impedance along each end-to-end channel on the PCBs and through the connector stages
- Planning signal pin assignments so that receive signal pins are grouped together while being isolated by a ground plane from the transmit signal pins (see Figure 10).
- Keeping signal traces short to minimize losses
- Routing signals in groups to minimize signal skew
- Reducing the number of through-hole via stubs by carefully selecting the layers to route the traces, controlling the PCB thickness, and back-drilling long via-hole stubs to minimize signal reflections

[4] Aggregate backplane bandwidth calculatian: 160 Gb/s x 16 blades x 2 directians = 5.12 Terabits/s

Downloaded from www.Manualslib.com manuals search engine

Figure 10. Separation of the transmit and receive signal pins by a ground plane in the in c-Class enclosure midplane

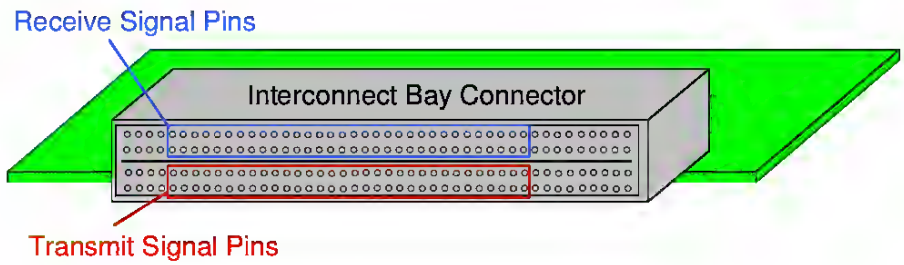

Separate power backplane

Distributing power on the same PCB that includes the signal traces would have greatly increased the board's complexity. Separating the power backplane from the NonStop signal midplane improves the signal midplane by reducing its PCB thickness, reducing electrical noise (from the power components) that would affect high-speed signals, and improving the thermal characteristics. These design choices result in reduced cost, improved performance, and improved reliability.

Channel topology and emphasis settings

Even when using best practices, high-speed signals transmitted across multiple connectors and long PCB traces can significantly degrade due to insertion and reflection losses. Insertion losses, such as conductor and dielectric material losses, increase at higher frequencies. Reflection losses are due to impedance discontinuities, primarily at connector stages. To compensate for these losses, a transmitter's signal waveform can be shaped by selecting the signal emphasis settings. However, the emphasis settings of a transmitter can depend on the end-to-end channel topology as well as the type of component sending the signal. Both of these can vary in the BladeSystem c-Class because of the flexible architecture and the use of mezzanine cards and embedded I/O devices such as network interface controllers (NICs). As shown in Figure 11, the topology for Device 1 on server blade 1 (a-b-c) is completely different than the topology for device 1 on server blade 4 (a-d-e). Therefore, an electronic keying mechanism in the Onboard Administrator identifies the channel topology for each device and ensures that the proper emphasis settings are configured for that device.

Downloaded from www.Manualslib.com manuals search engine

Figure 11. Different topologies require different emphasis settings

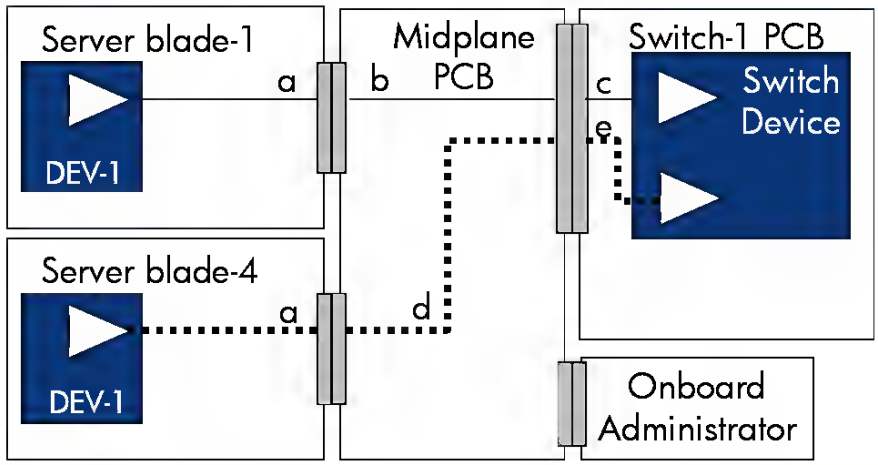

Signal midplane provides reliability

Finally, to provide high reliability, the NonStop signal midplane is designed as a completely passive board, meaning that it has no active components along the high-speed signal paths. The PCB consists primarily of traces and connectors. While there are a few components on the PCB, they are limited to passive devices that are extremely unlikely to fail. The only active device is an Electrically Erasable Programmable Read-Only Memory (EEPROM), which the Onboard Administrator uses to acquire information such as the midplane serial number. If this device were to fail, it would not affect the signaling functionality of the NonStop signal midplane. The NonStop signal midplane incorporates best design practices and is based on the same type of midplane used for decades in high-availability solutions such as the HP NonStop S-series, core networking switches from Cisco, Juniper Networks and core SAN switches from Cisco and Brocade. HP engineers have estimated that the mean time between failure (MTBF) for the signal midplane is in the hundreds of years.

Downloaded from www.Manualslib.com manuals search engine

Power backplane scalability and reliability

The power backplane is constructed of solid copper plates and integrated power delivery pins to ensure power distribution with minimum losses (Figure 12). Using solid copper plates reduces voltage drops, provides high current density, and high reliability.

Figure 12. Sketch of the c-Class power backplane showing the power delivery pins

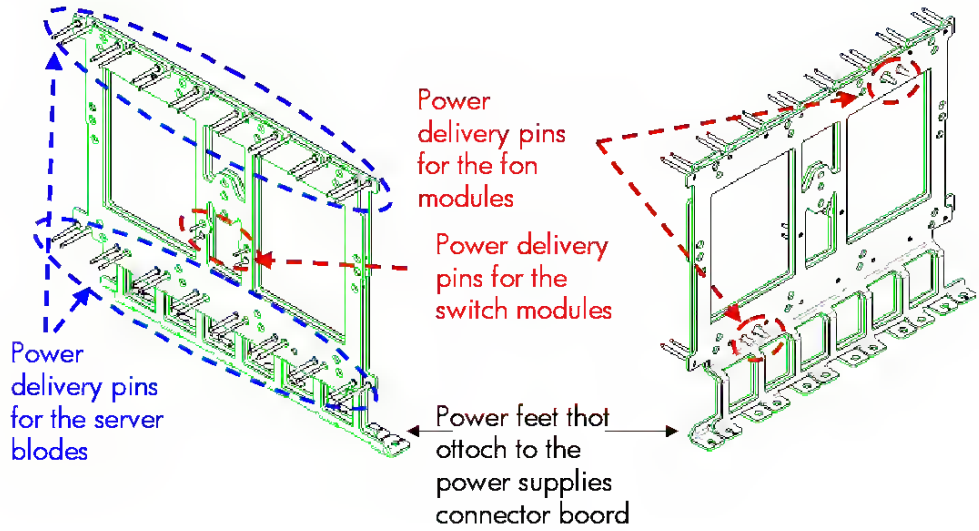
Power delivery pins for the fan modules
Power delivery pins for the switch modules
Power delivery pins for the server blades
Power feet that attach to the power supplies connector board

Power and cooling architecture with HP Thermal Logic

Power conservation and efficient cooling were key design goals for the BladeSystem c-Class. To achieve these goals, HP consolidated power and cooling resources, while efficiently sharing and managing them within the enclosure. HP uses the term *Thermal Logic* to refer to the mechanical features and control capabilities throughout the BladeSystem c-Class that enable IT administrators to optimize their power and thermal environments.

Thermal Logic encompasses technologies at every level of the c-Class architecture: processors, server blades, Active Cool fans, and the c-Class enclosure. Through the Onboard Administrator controller, IT administrators can access real-time power and temperature data, allowing them to understand their current power and cooling environments. Onboard Administrator allocates power to the device bays based on the specific configuration of each blade in the enclosure. As blades are inserted into the enclosure, the Onboard Administrator discovers each blade and allocates power accordingly, based on actual measured power requirements.

Onboard Administrator also allows customers to dynamically and automatically adjust operating conditions to meet their data center requirements. This allows them to maximize performance based on their power and cooling budgets and to forestall expensive power and cooling upgrades.

The technology briefs titled "HP BladeSystem c-Class c7000 enclosure technologies" and "HP BladeSystem c-Class c3000 enclosure technologies" give additional information about HP Thermal Logic technologies. Both are available on the HP technology website at www.hp.com/servers/technology.

Downloaded from www.Manualslib.com manuals search engine

Server blades and processors

At the processor level, HP Power Regulator for ProLiant[5] is a ROM-based power management feature of HP ProLiant servers. Power Regulator technology takes advantage of the power states available on x86 processors to scale back the power to a processor when it is not needed. Because the c-Class architecture shares power among all server blades in an enclosure, administrators can use Power Regulator technology to balance power loads among the server blades. As processor technology progresses, HP can recommend that customers use lower-power processor and component options when and where possible.

The server blade designs use precise ducting throughout the server blade to manage airflow and temperature based on the unique thermal requirements of all the critical components. The airflow is tightly ducted to ensure that no air bypasses the server blade and to obtain the most thermal work from the least amount of air. This concept allows much more flexibility in heat sink design choice. The heat sinks closely match the requirements of the server blade and processor architecture. For example, in the Intel® Xeon® based HP BladeSystem BL460c server blade, HP was able to use a smaller high-power processor heat sink than in rack-mount servers. These heat sinks have vapor chamber bases, thinner fins, and tighter fin pitch than previous designs. The smaller heat sink allows more space for full-size memory modules and hot plug hard drives on the server blades.

Most importantly, c-Class server blades incorporate intelligent management processors (Integrated Lights-Out 2, or iLO 2, for ProLiant server blades, or Integrity iLO for Integrity server blades) that provide detailed thermal information for every server blade. This information is forwarded to the Onboard Administrator and is accessible through the Onboard Administrator web interface.

Enclosure

At the enclosure level, HP Thermal Logic provides a number of advantages:

- Power designed to meet data center configurations
- High-efficiency voltage conversions
- Dynamic Power Saver mode to operate power supplies at high efficiencies
- Active Cool Fans that minimize power consumption
- Mechanical design features (PARSEC architecture) to optimize airflow

Meeting data center configurations

Rather than design the power budgets for the c-Class architecture based on the anticipated requirements of server blades, HP designed the c-Class enclosures to conform to typical data center facility power feeds. Thus, the enclosures are sized not only to amortize the cost of server blades across the infrastructure, but also to support the most server blades possible while using the power available today. As IT facilities managers choose to increase the number of power feeds into their facilities, c-Class enclosures can be added that will fit into those typical power feed budgets. Because the enclosures are sized to meets today's power infrastructure, there is no need for a separate power enclosure.

High-efficiency voltage conversions

Incorporating the power supplies into the enclosure reduced the distance over which power would need to be distributed. This allowed HP to use an industry-standard 12V infrastructure for the c-Class BladeSystem. Using a 12V infrastructure eliminates several power-related components and improves power efficiency on the server blades and infrastructure.

---

[5] For additional information about Power Regulator for ProLiant and which servers support it, see www.hp.com/servers/power-regulator.

Downloaded from www.Manualslib.com manuals search engine

Dynamic Power Saver Mode

Mast pawer supplies aperate inefficiently when lightly laaded and more efficiently when heavily laaded. When enabled, Dynamic Pawer Savings made will save power by running the required pawer supplies at a higher rate of utilization and putting unneeded power supplies in a standby made. When pawer demand increases, the standby power supplies instantaneously deliver the required extra power. As a result, the enclosure can operate at optimum efficiency, with no impact on redundancy. Both efficiency and redundancy are possible because the power supplies are consolidated and shared across the enclosure.

Active Cool fans

Quite often, small form-factor servers such as blade or 1U servers use very small fans designed to provide localized cooling in specific areas. Because such fans generate fairly low flow (in cubic feet per minute, or CFM) at medium back pressure, a single server often requires multiple fans to ensure adequate cooling. Therefore, when many server blades, each with several fans, are housed together in an enclosure, there is a trade-off between powering the fans and cooling the server blades. While this type of fan has proven to scale well in the BladeSystem p-Class, HP believed that a new design could better balance the trade-off between power and cooling.

A second solution for cooling is to use larger, blower-style fans that can provide cooling across an entire enclosure. Such fans are good at generating CFM, but typically also require higher power input, produce more noise, and must be designed for the highest load in an enclosure. Because these large fans cool an entire enclosure, failure of a single fan can leave the enclosure at risk of overheating before the fan is replaced.

With these two opposing solutions in mind, HP solved these problems by designing the Active Cool fan and by aggregating the fans to provide redundant cooling across the entire enclosure.

The Active Cool fans are controlled by the Onboard Administrator so that cooling capacity can be ramped up or down based on the needs of the entire system. Along with optimizing airflow, this control algorithm allows the c-Class BladeSystem to optimize acoustic levels and power consumption. Because of the mechanical design and the control algorithm, Active Cool fans deliver better performance—at least three times better than the next best fan in the server industry. As a result of the Active Cool fan design, the c-Class enclosures support full-featured servers that are 60 percent more dense than traditional rack-mount servers. Moreover, the Active Cool fans consume only 50 percent of the power typically required and use 30 percent less airflow. By aggregating the cooling capabilities of a few, high-performance fans, HP was able to reduce the overhead of having many, localized fans for each server blade, thereby simplifying and reducing the cost of the entire architecture.

PARSEC architecture

Each c-Class enclosure uses *PARSEC* (parallel, redundant, scalable, enclosure-based cooling) architecture. In this context, parallel means that fresh, cool air flows over all the server blades (in front of enclosure) and all the interconnect modules (in the back of the enclosure). Fresh air is pulled into the interconnect bays through a dedicated side slot in the front of the enclosure. Ducts move the air from the front to the rear of the enclosure, where it is then pulled into the interconnect modules and the central plenum, and then exhausted out the rear of the system.

Each power supply module has its own fan, optimized for the airflow characteristics of the power supplies. Because the power supplies and facility power connections are in a separate region of the enclosure, the fans can provide fresh, cool air and clear exhaust paths for the power supply modules without interfering with the airflow path of the server blades and interconnect modules.

Downloaded from www.Manualslib.com manuals search engine

Because the enclosures are designed to include separate physical cooling zones, the Active Cool fans provide cooling for their own zone and redundant cooling for the rest of the enclosure. One or more fans can fail and still leave enough fans to adequately cool the enclosure.[6]

To ensure scalability, HP designed both the fans and the power supplies with enough capacity to meet the needs of compute, storage, and I/O components well into the future.

HP optimizes the cooling capacity across the entire enclosure by optimizing airflow and minimizing leakage through the use of a relatively airtight central plenum, self-sealing louvers surrounding the fans, and automatic shut-off doors surrounding the device bays.

Configuration and management technologies

One of the goals of the BladeSystem c-Class architecture was to dramatically reduce the amount of time that IT personnel must spend to deploy and manage new systems. To achieve this design goal, HP provided an intelligent infrastructure that makes essential power and cooling information available to administrators and helps automate infrastructure management. Implementing the c-Class Onboard Administrator is the result of these goals.

The BladeSystem c-Class architecture also reduces the complexities of switch management in a blade environment. Blade environments provide distinct advantages because their direct backplane connections between switches and blades reduce the number of cables, and therefore cost and complexity. However, early generations of blade environments created the challenge of managing many additional small switches. HP solved this in an innovative way by developing Virtual Connect technology. Virtual Connect technology provides a way to virtualize the server I/O connections to Ethernet or Fibre Channel networks.

The technology briefs titled "Managing the HP BladeSystem c-Class" and "HP Virtual Connect technology implementation for the HP BladeSystem c-Class" provide detailed information about these technologies. They are available on the HP technology website: www.hp.com/servers/technology.

Integrated Lights-out technology

Each ProLiant server blade designed for the BladeSystem c-Class includes an iLO 2 management processor. The iLO 2 processor monitors thermal and operational conditions within each server blade and forwards this information on to the Onboard Administrator. Regardless of a server blade's operating condition, the iLO 2 management processor enables the remote management capabilities that customers have come to expect from ProLiant servers: access to a remote console, virtual media access, virtual power button, and system management information such as hardware health, event logs, and configuration. The iLO 2 device provides a higher-performance remote console (virtual KVM) as well as virtual media functionality that administrators can access from a web browser, command line, or script. The virtual KVM uses an architecture that acquires video directly from the video controller and uses an enhanced compression and refresh technology that reduces the amount of traffic on the network (thereby improving network efficiency).

Onboard Administrator

Onboard Administrator is a management controller module that resides within the BladeSystem c-Class enclosure. The Onboard Administrator controller communicates with the iLO 2 management processors on each server blade to form the core of the management architecture for BladeSystem c-Class. Customers have the option of installing a second Onboard Administrator board in the c7000

[6] The number of fans that can fail depends upon the number of blades, the number of fans, and the location of the blades. The OA reports thermal subsystem status and redundancy level, and will provide alerts when the thermal subsystem status changes.

Downloaded from www.Manualslib.com manuals search engine

enclosure to act as a redundant controller in an active-standby mode. In the future, HP plans to offer redundant Onboard Administrator capabilities for the c3000 enclosure.

Onboard Administrator collects system parameters related to thermal and power status, system configuration, and managed network configuration. It manages these variables cohesively and intelligently so that IT personnel can configure the BladeSystem c-Class and manage it in a fraction of the time that other solutions require.

Onboard Administrator monitors

- Thermal conditions. If the thermal load increases, the Onboard Administrator's thermal logic feature instructs the fan controllers to increase fan speeds to accommodate the additional demand.
- Power allocations guidelines and power capacity limits of various components. It uses sophisticated power measurement sensors to accurately determine how much power is being consumed and how much power is available. Because Onboard Administrator uses real-time, measured power data instead of maximum power envelopes, customers can deploy as many servers and interconnects as possible for the available power.
- Hardware configurations. One of the advantages of the BladeSystem c-Class is its flexibility in allowing customers to configure the system in virtually any way they desire. To assist in the configuration and setup process for the IT administrator, the Onboard Administrator verifies four attributes for each server blade and interconnect module as they are added to the enclosure: electronic keying of interconnects and mezzanine cards, power capacity, cooling capacity, and location of components. The electronic keying mechanism ensures that the interconnect modules and mezzanine cards are compatible. It also determines the signal topology and sets appropriate emphasis levels on the transmitters to ensure best signal reception by the receiver after the signal passes across the high-speed NonStop signal midplane.
- Management network control capabilities. Onboard Administrator provides tools to automatically identify and assign IP addresses for the BladeSystem c-Class components on existing management networks (for components supporting Dynamic Host Configuration Protocol). This simplifies and automates the process of configuring the BladeSystem c-Class.

A simple Insight Display screen on the front of each HP BladeSystem c-Class enclosure (Figure 13) provides quick, on-site access to all the setup, management, and troubleshooting features of the Onboard Administrator. For example, when the enclosure is powered up for the first time, the Insight Display launches an installation wizard to guide an IT technician through the configuration process. After the technician initially configures the enclosure, the Insight Display provides feedback and advice if there are any installation or configuration errors. In addition, the technician can access menus that provide information about Enclosure Management, Power management, and HP BladeSystem Diagnostics. The Insight Display provides a User Note function that is the electronic equivalent of a *sticky note*. Administrators can use this function to display helpful information such as contact phone numbers or other important information. The Insight Display also provides a bi-directional chat mode (similar to instant messaging) between the Insight Display user (such as a service technician) and a remote administrator using the web GUI. This allows a technician in the data center to communicate instantly with a remote administrator about what needs to be done.

Downloaded from www.Manualslib.com manuals search engine

Figure 13. The main menu an the Insight Display

IT technicians and administrators can access the Onboard Administrator in any of three ways: through the Insight Display, through a web graphical user interface (GUI), or through a command-line interface (CLI). A fourth access method is available for the c3000 enclosure: The optional c3000 KVM Module provides the Onboard Administrator CLI, Insight Display, and KVM console connections to all the server blades in the enclosure.

Virtuolized network infrostructure with Virtuol Connect technology

HP BladeSystem c-Class is designed from the ground up integrating Virtual Connect technology. The OnBoard Administrator, the c-Class PCI-Express mezzanine cards, the embedded NICs, and iLO 2 all provide functionality to support the Virtual Connect technology. The tight integration of the Virtual Connect capability into the HP BladeSystem c-Class infrastructure makes its functionality highly effective and seamless.

Virtual Connect implements server-edge virtualization: It puts an abstraction, or virtualization, layer between the servers and the external networks so that the local area network (LAN) and storage area network (SAN) see a pool of servers rather than an individual server (see Figure 14) behind a physical port. Specific interconnect modules—Virtual Connect modules—provide the virtualized connections so that a server can be replaced or a workload moved from one server to another without disrupting the LAN and SAN configurations. This enables the system administrator to be self-sufficient – to move workloads, add servers, or replace servers without needing help from the network and storage administrators.

Once the LAN and SAN connections are made to the Virtual Connect modules to pool the servers, the system administrator uses a Virtual Connect Manager User Interface to create a server identity. The Virtual Connect Manager (VC Manager) creates a server-specific I/O profile, consisting of unique media access control (MAC) addresses for all NICs and World Wide Names (WWN) for all host bus adapters (HBAs). Then, the VC Manager replaces each server's default MAC and WWN addresses with the addresses created in the server-specific I/O profile. These profiles are managed locally. If any changes need to occur (for instance, if a server blade needs to be upgraded), the system administrator can swap out the server blade, and the Virtual Connect Manager will manage the NIC and HBA address changes.

Downloaded from www.Manualslib.com manuals search engine

Figure 14. HP Virtual Connect technology provides server-edge virtualization

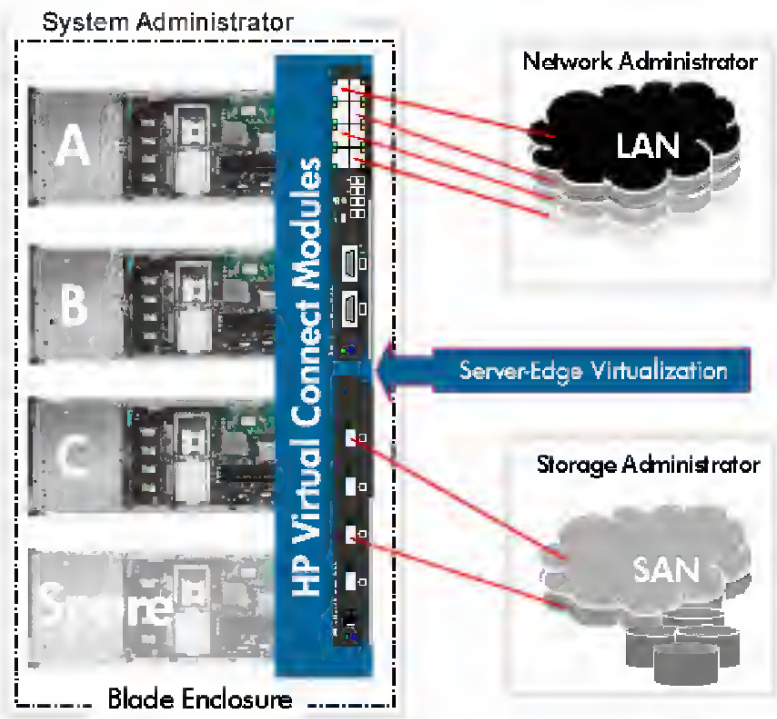

From the perspective of the network and storage administrators, the LAN and SAN connections are established with the pool of servers, and the administrators see no changes to their networks. This allows moving workloads and adding or replacing servers without affecting the LAN or SAN. In addition, the Virtual Connect modules do not participate in network control activities (such as Spanning-Tree Protocol for Ethernet or FSPF for Fibre Channel) as a switch would. Therefore network administrators will not have extra switches to manage on the edge of their networks.

HP Virtual Connect technology provides a simple, easy-to-use tool for managing the connections between HP BladeSystem c-Class servers and external networks. It cleanly separates server management from LAN and SAN management, relieving LAN and SAN administrators from server maintenance.

More information about HP Virtual Connect technology is available in the technology brief titled "HP Virtual Connect technology implementation for the HP BladeSystem c-Class" available at www.hp.com/servers/technology.

HP Virtual Connect Enterprise Manager is a software application that simplifies the management of large BladeSystem environments using Virtual Connect to control LAN and SAN connectivity. It allows administrators to seamlessly manage multiple c-Class enclosures using a single console application. The Virtual Connect Enterprise Manager provides a central pool of Virtual Connect LAN and SAN addresses for multiple enclosures. It allows administrators to logically link separate Virtual Connect domains and move server profiles between those domains, as long as the servers are physically connected to the same network. The single resource pool and capability to group large numbers of servers enables rapid, reliable deployment and movement of servers across the datacenter.

More information about Virtual Connect Enterprise Manager is available at www.hp.com/go/vcem.

Downloaded from www.Manualslib.com manuals search engine

Availability technologies

The BladeSystem c-Class incarparates layers af availability ta enable the 24/7 infrastructure needed in data centers. Using redundant canfiguratians eliminates single paints af failure, and the c-Class architecture reduces the risk af component failures and the time required for changes.

Redundont configurotions

The BladeSystem c-Class minimizes the chances af a failure by praviding redundant pawer supplies, fans, and intercannect modules. For example, custamers have the aptian af using power supplies in an N+N redundant canfiguration or an N+1 canfiguratian. The intercannect madules can be placed side-by-side far redundancy, as shawn in Figure 6 on page 10. And the c7000 enclasure is capable of supparting either one ar two Onboard Administrator modules in an active-standby configuration. Redundant Onboard Administrator functionality is also planned for the c3000 enclosure.

The c-Class architecture provides redundant paths through the use of multiple facility power feeds into the enclosures, blade-to-interconnect bay connectivity, and blade-to-enclosure manager connectivity. Because all c-Class components are hot-pluggable, administrators can quickly re-establish a redundant configuration in the event of a failure.

Relioble components

HP took every opportunity in the c-Class architecture to design for reliability, especially for critical components that can be considered single points of failure. Some customers might consider the NonStop signal midplane for the BladeSystem c-Class enclosure to be a single point of failure, since it is not replicated. However, HP mitigated this risk and made the PCB extremely reliable:

- Designed the NonStop signal midplane to provide redundant paths between the server blades and interconnect bays
- Eliminated all active components from the PCB that would affect functionality, thereby removing potential sources of failure
- Removed power from the NonStop signal midplane to reduce board thickness, reduce thermal stresses, and reduce the risk of any power bus overloads affecting the data signals
- Minimized the connector count to reduce mechanical alignment issues
- Used mechanically robust midplane connectors that also support at least 10 Gigabit per second (Gbps) high-speed signals with minimum crosstalk

The result is a NonStop signal midplane that has a calculated MTBF in the hundreds of years.

Some c7000 customers may choose to have a single Onboard Administrator module rather than two for redundancy. In this case, the Onboard Administrator can be a single point of failure. In the unlikely event of an Onboard Administrator failure, server blades and interconnect modules will all continue to operate normally. The module can be removed and replaced without affecting operations of the server blades and interconnect modules.

Operating temperatures of components can play a significant role in reliability. As the operating temperature increases beyond specified maximum values, thermal stresses increase, which results in shortened life spans. The PARSEC architecture of the BladeSystem c-Class enclosures minimizes the operating temperature of components by delivering fresh, cool air to all critical components. The airflow is tightly ducted to make every cubic centimeter of airflow count and to ensure the most thermal work from the least amount of air. The server blades are designed with ample room for intake air and heat sinks (both on the processor and memory modules). Rather than use the traditional heat sink design for the processors, HP designed a copper-finned heat sink that provides more heat transfer in a smaller package than traditional heat sinks used in 1U rack-optimized servers.

Downloaded from www.Manualslib.com manuals search engine

Finally, Onboard Administrator's thermal monitoring of the entire system ensures that the Active Cool fans deliver adequate cooling to the entire enclosure. Because the fan design uses a high-performance motor and impeller, it consumes less power and uses less airflow to cool an enclosure than a traditional fan design would. Because of its unique fan blade, housing, motor windings, bearings, and drive circuit, the Active Cool fan provides higher reliability than typical server fans.

Reduced logistical delay time

Like other ProLiant servers, the BladeSystem c-Class was designed with ease-of-use as a priority. Several important technologies in the BladeSystem c-Class reduce the amount of time needed to replace, upgrade, and configure systems: Onboard Administrator and the related Insight Display, Virtual Connect technology, and hot-plug devices.

Onboard Administrator and Virtual Connect technologies have already been discussed in the section titled "Configuration and management technologies." With the intelligence of the Onboard Administrator and the easy-to-use Insight Display panel, administrators can configure and troubleshoot their systems in minutes, rather than hours or days. Adopting Virtual Connect technology removes administrative burdens from LAN and SAN administrators because that they are not required to change their network setup every time a configuration change occurs within the server blade environment. Furthermore, because the network connections are made to a pool of server blades, it is quick and easy to migrate the network service from a failed server blade to a functional server blade.

Finally, the fans, power supplies, interconnect modules, Onboard Administrator modules, server blades, and storage blades are hot-pluggable, meaning that they can be removed without affecting any other components in the enclosure.

Conclusion

HP designed the BladeSystem c-Class as an architecture that would deliver on the promise of a modular, adaptive, automated data center. To do this, HP worked very closely with its customers to understand their requirements and challenges in managing their data centers. By combining this knowledge with the recognition of emerging industry standards and technologies, architects from multiple business units within HP collaborated to define the c-Class architecture, enclosure design, and Thermal Logic cooling technologies.

The c-Class architectural model provides scalable device and interconnect bays plus power and cooling headroom that allows customers to add the components they need, when they need them. Customers can easily scale the enclosure from the minimum of one server blade to the maximum by adding more fans and power supplies. By designing a unique NonStop signal midplane that can adapt to customer needs and technology directions over multiple generations, HP has ensured flexibility and a long life for the BladeSystem c-Class. By consolidating key resources—volume space, power, cooling, and signal traces across the midplane—the BladeSystem c-Class ensures that resources can be shared efficiently to meet customer needs. The c-Class architecture is designed for longevity and interoperability with server blades, storage blades, and interconnect modules for several generations of products. The c-Class architecture enables multiple enclosure designs, optimized for different market segments.

With the BladeSystem c-Class, HP has delivered even more hardware control, intelligent monitoring, automation capabilities, and virtualization capabilities than with previous generations of blade systems. The Onboard Administrator and Insight Display work in conjunction with the intelligent management processors on each server blade to provide information and control to administrators. In addition, HP has differentiated the BladeSystem c-Class from its competitors through HP Thermal Logic that dynamically monitors and controls power and cooling in an extremely cost-effective manner, and HP Virtual Connect technology that simplifies network management and IT changes by virtualizing I/O connections.

Downloaded from www.Manualslib.com manuals search engine

For more information

For additional information, refer to the resources listed below.

Source	Hyperlink
HP BladeSystem c-Class documentation	http://h71028.www7.hp.com/enterprise/cache/316735-0-0-0-121.html
HP BladeSystem power sizer	www.hp.com/go/bladesystem/powercalculator
HP BladeSystem website	www.hp.com/go/bladesystem/
HP Power Regulator for ProLiant	http://h18000.www1.hp.com/products/servers/management/ilo/power-regulator.html
HP Technology Briefs: ▪ HP BladeSystem c-Class c7000 enclosure technologies ▪ HP BladeSystem c3000 Enclosure technologies ▪ HP Virtual Connect technology implementation for the HP BladeSystem c-Class ▪ Managing the HP BladeSystem c-Class ▪ HP ProLiant c-Class Server Blades ▪ HP BladeSystem c-Class SAN connectivity	http://h18013.www1.hp.com/products/servers/technology/whitepapers/proliant-servers.html#bl
HP Labs technical report: Blades as a General-Purpose Infrastructure for Future System Architectures: Challenges and Solutions	www.hpl.hp.com/techreports/2006/HPL-2006-182.pdf

Downloaded from www.Manualslib.com manuals search engine

Call to action

Send comments about this poper to TechCom@HP.com.

© Copyright 2006, 2008 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Intel and Intel Xeon are trademarks of Intel Corporation in the United States and other countries.

TC080301T8, March 2008

Downloaded from www.Manualslib.com manuals search engine